CSRF Interim Project Report: Current Topics in Density Functional Theory

Richard P. Muller, Marcus G. Martin, Peter Schultz, and Renee M. Van Ginhoven

Multiscale Computational Materials Models (1435), Sandia National Laboratories, Albuquerque, NM, 87185-1110

This report describes progress made under the CSRF project *Current Topics in Density Functional Theory*. Density functional theory (DFT) is a technique for computing the energetics of molecules and materials. DFT is unique among atomistic simulation techniques in combining a high accuracy approximation to the quantum mechanics describing the chemical bonding in materials with computationally tractible solutions. This combination makes it particularly important and relevant for Sandia's stockpile stewardship mission, a large part of which involves insuring material performance without nuclear testing. DFT techniques have already made important contributions to Sandia's DP programs. However, to have greater and wider impact at Sandia, DFT techniques need to have **increased accuracy**, they need to provide this accuracy with **increased speed**, and they need to have **increased range** as to the problems and conditions they can contribute to. The goal of this project is to achieve these improvements to DFT.

Contents

I. INTRODUCTION

This report describes progress made under the CSRF project *Current Topics in Density Functional Theory*. Density functional theory (DFT) is a quantum mechanical technique for computing the energetics of molecules and materials. DFT is unique among atomistic simulation techniques in combining a high accuracy approximation to the quantum mechanics describing the chemical bonding in materials with computationally tractible solutions. This combination makes it particularly important and relevant for Sandia's stockpile stewardship mission, a large part of which involves insuring future material performance without nuclear testing. DFT techniques have already made important contributions to Sandia's DP programs. However, to have greater and wider impact at Sandia, DFT techniques need to have **increased accuracy**, they need to provide this accuracy with **increased speed**, and they need to have **increased range** as to the problems and conditions they can contribute to. The goal of this project is to achieve these improvements to DFT.

This report will detail progress made in each of these areas. We will begin by describing the accuracy of existing density functionals and how the *Exact Exchange* functionals can improve upon this. We will then discuss our attempts to improve the speed of DFT with density matrix-based methods that replace the matrix eigenproblem with techniques that determine the density matrix directly from the electronic Hamiltonian. Finally, we will describe our efforts to broaden the range of problems that may be treated with DFT, focusing particularly on the implementation of configurational biased Monte Carlo techniques for liquid densities and liquid-vapor coexistance curves.

II. OVERVIEW OF DENSITY FUNCTIONAL THEORY

Density functional theory (DFT) dates from the work of Thomas, Fermi, and Dirac (1–3) describing the ground electronic state of an atom, molecule, or material in terms of its electronic density $\rho(\mathbf{r})$

$$E[\rho] = T[\rho] + v_n[\rho] + v_J[\rho] + v_x[\rho] + v_c[\rho], \qquad (1)$$

where T is the kinetic energy of an electron, v_n is the nuclear attraction each electron feels, v_J is the electron-electron repulsion, v_x is the "exchange" interaction that insures the fluid of electrons act as a Fermi-Dirac fluid, and v_c is the correlation energy. The determination of the kinetic energy functional was a major challenge for the development of DFT approaches. In 1965, Kohn and Sham (4) re-introduced one-particle orbital eigenfunctions to DFT as a way of more accurately computing the kinetic energy functional, and the accuracy of the resulting calculations have been such that these calcula-

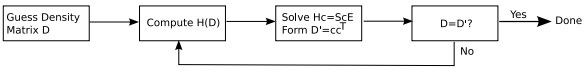

FIG. 1 Flow chart of where time is spent in a DFT calculation.

tions have dominated DFT calculations ever since. The resulting equations have the form

$$(t + v_{\text{eff}}[\rho])\phi_i = \epsilon_i \phi_i, \tag{2}$$

where t is the kinetic energy operator,

$$t\phi_i = -\frac{\hbar^2}{2m_e}\nabla^2\phi_i, \tag{3}$$

and v_{eff} is the effective potential,

$$v_{\text{eff}}[\rho] = v_n + v_J[\rho] + v_x[\rho] + v_c[\rho], \tag{4}$$

composed of the nuclear attraction potential v_n, the electron repulsion potential v_J, the exchange potential v_x, which insures that the density behaves according to Fermi-Dirac statistics, and the correlation potential v_c, which subsumes all of the remaining many-body interactions. These equations can be written in the form of a matrix eigenvalue equation

$$\mathbf{HC} = \mathbf{C}\Lambda \tag{5}$$

$$\mathbf{C}^T\mathbf{C} = \mathbf{I}. \tag{6}$$

The columns of $\mathbf{C}$ and diagonal elements of Λ are the orbital eigenvectors describing the electronic wave function and corresponding orbital energies for the first N_{occ} (typically $N_{occ} = N_{el}/2$) occupied orbitals. We will find it convenient to define the (one-particle) *density matrix* $\mathbf{D}$ as

$$\mathbf{D} = \mathbf{C}\mathbf{C}^T \tag{7}$$

Figure 1 shows a diagram of the computational steps that comprise a DFT calculation. There are two time-consuming steps. There is a *Hamiltonian formation step* that forms the Kohn-Sham Hamiltonian operator for a given system for a specified density matrix $\mathbf{D}$, and there is a *Hamiltonian solution step* that solves the eigenproblem associated with that Hamiltonian and forms an updated $\mathbf{D}$ matrix; these two steps are iterated until self-consistency is reached. Formally, the Hamiltonian formation step scales $\mathcal{O}(N^4)$ for a system of size N. However, intelligent Hamiltonian formation steps have reduced the scaling of this step to roughly $\mathcal{O}(N)$ for large systems. The eigenproblem formally scales $\mathcal{O}(N^3)$, and this scaling now dominates the calculation time for large systems.

Figure 2 shows the timings of the Quest DFT program. The red and the blue lines show the timings for Hamiltonian formation during the setup phase and during each iteration, respectively. The green line shows the timings for the solution of the eigenproblem. Beyond roughly 150 atoms, the eigensolve dominates the problem.

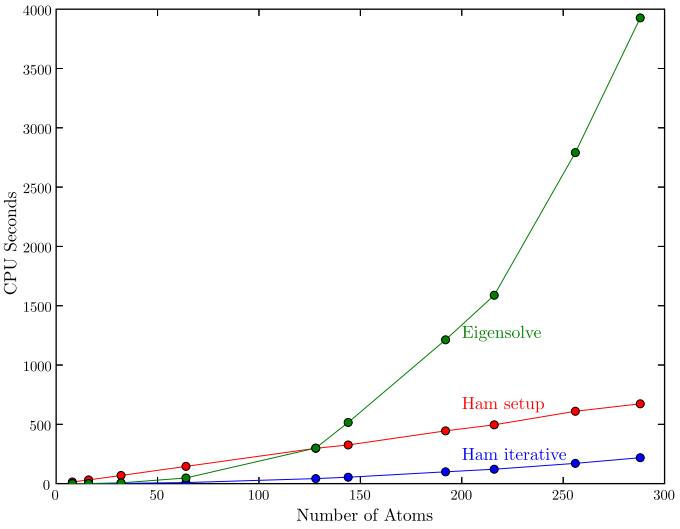

FIG. 2 Timings for the Quest DFT program for Si crystals with differing number of atoms. The red and the blue lines show the timings for Hamiltonian formation during the setup phase and during each iteration, respectively. The green line shows the timings for the solution of the eigenproblem.

The first DFT functionals with qualitative accuracy arose from Ceperley and Alder's Quantum Monte Carlo (QMC) calculations for the electron gas (5), which were fitted to DFT correlation functionals by Vosko *et al.* (6), using a Slater functional (7) interaction. This functional used *Local Density Approximation* (LDA) to write the exchange-correlation energy as a functional of the local electronic density at each point in space:

$$E_{xc}^{LDA} = \int F[\rho(r)]\rho(r)dr. \tag{8}$$

These functionals were later replaced by functionals using the *Generalized Gradient Approximation* (GGA) (8–10) that wrote the functionals in terms of not only the local density, but also the gradient of that density, at each point in space:

$$E_{xc}^{GGA} = \int F[\rho(r), \nabla\rho(r)]\rho(r)dr. \tag{9}$$

The GGA funtionals themselves were later replaced by *Hybrid Functionals* (11; 12) that took a linear combination of the exchange-correlation operator from GGA with the exact exchange matrix from Hatree-Fock (HF) theory. This combination provided greater accuracy, but it came at the loss of some theoretical rigor, as density functionals are functionals of the electronic density, whereas the HF exchange matrix is a functional of the molecular orbitals.

Table I shows the performance of different DFT functionals on the G2 themochemical test suite, a set of small molecules chosen by Pople *et al.* (13) for which

TABLE I A comparison of the errors on the G2 themochemical test suite of small molecules for different DFT functionals. Shown is the LDA functional SVWN, the GGA functionals BLYP, BP86, BPW91, and the hybrid functionals B3LYP, B3PW91, and B3P86. Data for the table from reference (13).

Functional	Year	Avg. Error	Max. Error
SVWN	1980	39.60	93.80
BLYP	1988	4.69	15.30
BP68	1986	10.50	26.10
BPW91	1991	5.20	19.10
B3LYP	1993	2.43	8.42
B3PW91	1993	2.59	7.40
B3P86	1993	7.84	24.60

high-quality themochemical data was available. The LDA functional shows an average accuracy of roughly 40 kcal/mol. The GGA functionals do substantially better, with an average accuracy of 5-10 kcal/mol. The best of the hybrid functionals have an average accuracy of roughly 2.5 kcal/mol.

In this discussion it is worth considering how much energy is sufficient to describe materials chemistry. A useful yardstick is that a $C-C$ chemical bond provides roughly 100 kcal/mol energy, and that a typical molecule will consist of several bonds of this strength. A second yardstick is that in computing the rates of chemical reactions, an error of 1.5 kcal/mol in computing an activation barrier leads to roughly an order of magnitude error in computing the rate. We will thus take the somewhat arbitrary figure of **1.5 kcal/mol** as our measure of chemical accuracy, as calculations that can produce this level of accuracy will have smaller than a 1% error for molecules with multiple bonds, and should lead to preditions of rates of reactions that are correct within an order of magnitude.

III. INCREASING THE ACCURACY OF DFT

A. Exact exchange density functionals

The goal of the exact exchange (EXX) method (14–18) is to obtain density functionals from orbital-dependent energy functions. These methods have allowed the generation of exchange functionals that derive from the HF exchange energy expression and yet are true density functionals. These methods can generate functionals that have the advantages provided by hybrid functionals, but with greater theoretical rigor.

There has been considerable hope that EXX methods will help solve the *band gap* problems that plague HF and DFT. In HF theory, occupied orbitals are properly self-interaction corrected, whereas unoccupied orbitals are not (19), leading to unnaturally high energies for the unoccupied orbitals, and a band gap that is too large. In contrast, in DFT both the occupied and unoccupied orbitals are only self-interaction corrected to the extent that the DFT exchange operator cancels the self-Coulomb interaction. This cancellation is typically imperfect, and DFT typically displays band gaps that are too small, although the fact that both the occupied and unoccupied orbitals display similar errors yields band gaps that are in general closer to the correct values than those from HF theory. The methods based upon EXX are rigorously self-interaction corrected, and, at the same time, have a consistent treatment of occupied and unoccupied states, giving hope that these approaches will combine the strengths of HF and DFT approaches. Preliminary studies on the computation of band gaps (18; 20; 21), and, in particular, those including a correlated orbital-dependent functional (22; 23), suggest that this hope is well-founded.

1. Exact exchange methodology

The KS equations take the form of eq (2). In the EXX method the potential v_{eff} is determined that minimizes some objective energy functional via the variations $\delta E/\delta v_{\text{eff}}$. This energy functional is typically taken as the HF energy, but in the current work we explore other choices for this functional as well.

The main idea behind EXX is to use *chain rule* type derivatives to enable proper density functionals to be constructed from orbital dependent energy functionals:

$$\frac{\delta E}{\delta v_{\text{eff}}} = \sum_i \frac{\delta E}{\delta \phi_i}\frac{\delta \phi_i}{\delta v_{\text{eff}}}. \tag{10}$$

We follow Yang and Wu (24), who define a particularly elegant technique whereby the EXX is expanded in a set of Gaussian functions $g(r)$ about a reference potential $v_0(r)$

$$v_{\text{eff}}(r) = v_n(r) + v_0(r) + \sum_\ell b_\ell g_\ell(r). \tag{11}$$

Yang and Wu take the reference potential $v_0(r)$ to be the Fermi-Amaldi potential

$$v_0(r) = \frac{N-1}{N}\int \frac{\rho_0(r')}{|r-r'|}dr', \tag{12}$$

which turns out to be a very good approximation to the proper EXX potential. Because the reference potential is independent of v_{eff}, the derivative may be obtained by minimizing the energy functional with respect to the expansion coefficients b_ℓ via

$$\begin{aligned}\frac{\delta E[\{\phi_i\}]}{\delta v_{\text{eff}}} &= \sum_\ell \frac{\delta E[\{\phi_i\}]}{\delta b_\ell} \\ &= \sum_\ell \sum_{i,a\neq i} \int \frac{\delta E[\{\phi_i\}]}{\delta \phi_i(r)}\phi_a(r)dr \frac{\langle \phi_a | g_\ell | \phi_i\rangle}{\epsilon_i - \epsilon_a}\end{aligned} \tag{13}$$

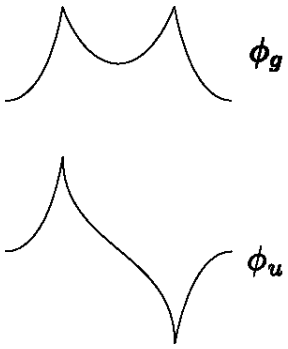

FIG. 3 The first and second natural orbitals for H_2.

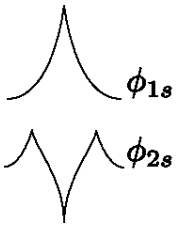

FIG. 4 The first and second natural orbitals for He.

In the special case of $E[\{\phi_i\}]$ equal to the HF total energy, this equation simplifies to (24)

$$\frac{\delta E^{EXX}[\{\phi_i\}]}{\delta v_{\text{eff}}} = \sum_{\ell} \sum_{i,a\neq i} \int \langle \phi_i \left| F^{HF} \right| \phi_a \rangle \frac{\langle \phi_a \left| g_\ell \right| \phi_i \rangle}{\epsilon_i - \epsilon_a} \tag{14}$$

where F^{HF} is the Hartree-Fock Fock operator.

We note, however, that the EXX method is not limited to using the HF energy functional. The next section describes a new breakthrough that has come from our work in using a multiconfigurational approach to introduce not only exact exchange but also major contributions from of correlation energy.

2. GVB and EXX-GVB

The HF wave function contains only a single electronic configuration: in the HF description of H_2 both electrons occupy the bonding orbital ϕ_g, shown in Figure 3, and in the HF description of He both electrons occupy the orbital ϕ_{1s}, shown in Figure 4. This single configuration limits the ways the wave function can be variationally minimized and leads to, for example, the well-known problem of the HF description of H_2 dissociating to the wrong limits, shown in Figure 5. It is possible to dissociate H_2 using an unrestricted HF (UHF) description, but the resulting wave function contains spin contamination and is thus no longer an eigenfunction of the spin S^2 operator.

The GVB-PP wave function (25) adds a variationally

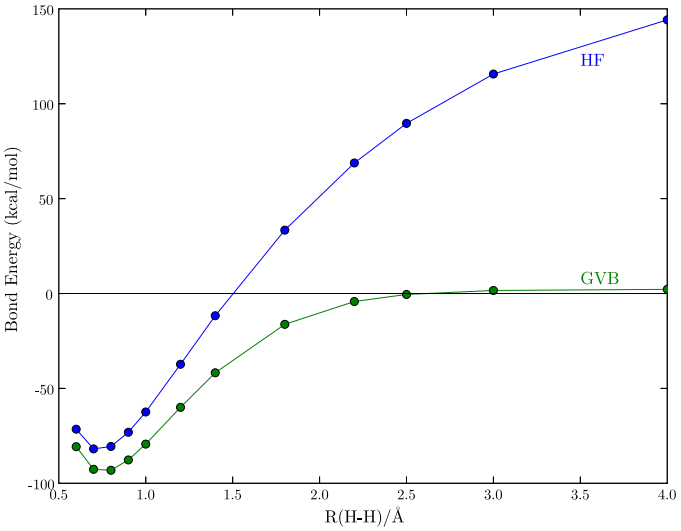

FIG. 5 A comparison of H_2 dissociation using the one-configuration HF wave function (blue circles) and the two-configuration GVB-PP wave function (green circles). Only the GVB wave function dissociates to the correct limit.

determined amount of another electronic configuration that gives the overall wave function greater freedom. In H_2, the GVB-PP wave functions adds the antibonding orbital ϕ_u, shown in Figure 3, which allows the two electrons to build in some static electron correlation and to dissociate properly, shown in Figure 5. In He, the GVB-PP wave function adds the second atomic orbital ϕ_{2s}, shown in Figure 4, which allows the two electrons to build in some static electron correlation and avoid each other to some degree. In the GVB terminology, the orbitals ϕ_g and ϕ_u, or ϕ_{1s} and ϕ_{2s}, are the first and second *natural orbitals* of a *GVB pair*.

For two electron systems with one GVB pair, the GVB-PP wave function takes a particularly simple form. For the purposes of our discussion, we will refer to the first and second natural orbitals of the GVB pair as ϕ_a and ϕ_b.

$$\Psi^{GVB} = (c_a\phi_a^2 - c_b\phi_b^2)\alpha\beta. \tag{15}$$

ϕ_a and ϕ_b are spatial orbitals, and α and β are the corresponding spin components. The coefficients c_a and c_b are determined by solving the 2×2 configuration interaction (CI) matrix:

$$\mathbf{HC} = \Lambda\mathbf{C} \tag{16}$$

where

$$\mathbf{H} = \begin{pmatrix} E_a & K_{ab} \\ K_{ab} & E_b \end{pmatrix}. \tag{17}$$

The CI matrix elements are given by

$$E_a = 2h_{aa} + (aa|aa) \tag{18}$$

$$E_b = 2h_{bb} + (bb|bb) \tag{19}$$

$$K_{ab} = (ab|ab) \tag{20}$$

The one-electron terms h_{ij} are given by

$$h_{ij} = \int \phi_i(r) h(r) \phi_j(r) dr \tag{21}$$

and h contains the kinetic-energy and nuclear attraction terms. The two-electron terms $(ij|kl)$ are in *chemist's notation* (see, for example, reference (26)) and are given by

$$(ij|kl) = \int \frac{\phi_i^*(r_1)\phi_j(r_1)\phi_k^*(r_2)\phi_l(r_2)}{|r_1 - r_2|} d^3r_1 d^3r_2. \tag{22}$$

After the CI coefficients are determined, the GVB-PP energy may be determined as

$$E^{GVB} = c_a^2 E_a + c_b^2 E_b + 2c_a c_b K_{ab} + E_{ZZ'} \tag{23}$$

where $E_{ZZ'}$ is the nuclear repulsion energy.

In the current work we take eq (23) to be the objective function to be optimized via eq (14). The derivatives $\delta E_{GVB}/\delta b_\ell$ are straightforward given the relations

$$\int \frac{\delta E^{GVB}}{\delta \phi_a(r)} \phi_j(r) dr = c_a^2 (h_{aj} + (aa|aj)) + c_b c_a (ba|bj), \tag{24}$$

$$\int \frac{\delta E^{GVB}}{\delta \phi_b(r)} \phi_j(r) dr = c_b^2 (h_{bj} + (bb|bj)) + c_b c_a (ba|ja). \tag{25}$$

3. Results on Band Gaps of He and Be

In KS DFT, only the highest occupied eigenvalue has a true physical interpretation, corresponding to the negative of the lowest ionization energy. In reference (27), Savin, Umrigar and Gonze derive a nearly exact Kohn-Sham potential from quantum Monte Carlo (QMC) calculations, and demonstrate that the resulting KS eigenvalues reproduce values from experiment (28) and explicit Hylleraas coordinate calculations of the excited states (29; 30). The experiments and Hylleraas calculations give different values for the singlet and triplet excited states; in contrast, the KS eigenvalues from the QMC exchange-correlation functional yield only a single, spin-averaged value for each state. The fact that these values fall between the singlet and triplet energies for each state is a remarkable result, which the authors interpret as evidence that the Kohn-Sham orbitals arising from their QMC-based Kohn-Sham potential and the exact quasiparticle orbitals obey the same long-range equations to order $1/r^4$. For the remainder of this paper we will take the KS eigenvalues from reference (27) to be the "correct" values.

In the current work we report results using an EXX functional as well as a new method (EXX-GVB) derived from the GVB-PP wave function. Our results do not display the same quantitative agreement with experiment and Hylleraas calculations that those in reference (27) do, our methods come with substantially less computational expense. The accuracy of our results demonstrates that this approach does, in fact, exhibit the correct long-range behavior, and provides hope that, with an appropriate choice of a correlation functional, inexpensive DFT calculations might yield the quantitative accuracy that Savin, Umrigar, and Gonze's QMC-based KS-DFT calculations provided.

Table II reports a comparison of Helium excitation energies to the QMC-derived exchange-correlation functional as well as to HF and to standard B3LYP functionals. We report the mean absolute deviation (MAD) between the QMC-derived values and those from HF and the DFTs. The HF excitation energies differ on the average by 0.107 h, the LDA, BLYP, and PBE values differ

TABLE II Comparison of the Helium excited state spectrum to betwen QMC-derived KS potential (reference (27)), and those from HF, LDA, BLYP, PBE, B3LYP, and our current EXX approach. The final line reports the mean absolute deviation (MAD) between the QMC-derived exchange-correlation functional and the other techniques. Energies are reported in Hartree atomic units.

State	QMC	HF	LDA	BLYP	PBE	B3LYP	EXX
1s→2s	0.746	0.918	0.571	0.585	0.580	0.662	0.762
1s→2p	0.777	0.923	0.576	0.590	0.585	0.667	0.793
1s→3s	0.839	0.922	0.574	0.587	0.582	0.665	0.856
1s→3p	0.848	0.941	0.592	0.605	0.600	0.683	0.864
1s→3d	0.848	0.936	0.589	0.603	0.598	0.680	0.865
1s→4s	0.869	0.929	0.580	0.594	0.591	0.673	0.885
MAD		0.107	0.241	0.227	0.232	0.150	0.016

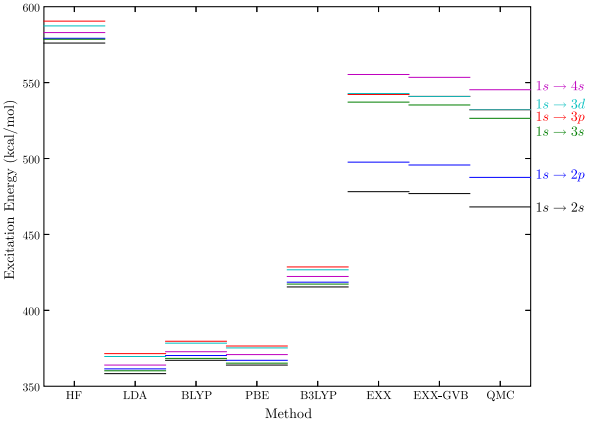

FIG. 6 Comparison of excitation energies for He atom computed from the KS-DFT eigenvalues using HF, LDA, BLYP, PBE, B3LYP, EXX, and QMC methods. The levels are color-coded based on the excitation, shown at the right. The OEP methods are the only ones that obtain the correct ordering of the states, and the errors are reasonably constant across the whole spectrum, in contrast with the other approaches.

TABLE III Comparison of the Beryllium excited state spectrum between QMC-derived KS potential (reference (27)), and those from HF, LDA, BLYP, PBE, B3LYP, and our current EXX approach. The final line reports the mean absolute deviation (MAD) between the QMC-derived exchange-correlation functional and the other techniques. Energies are reported in Hartree atomic units.

State	QMC	HF	LDA	BLYP	PBE	B3LYP	EXX
2s→2p	0.133	0.313	0.129	0.130	0.132	0.180	0.130
2s→3s	0.244	0.312	0.205	0.198	0.204	0.230	0.240
2s→3p	0.269	0.325	0.210	0.204	0.210	0.236	0.267
2s→3d	0.283	0.331	0.220	0.214	0.220	0.246	0.278
2s→4s	0.296	0.321	0.211	0.206	0.212	0.238	0.292
MAD		0.075	0.050	0.055	0.050	0.038	0.004

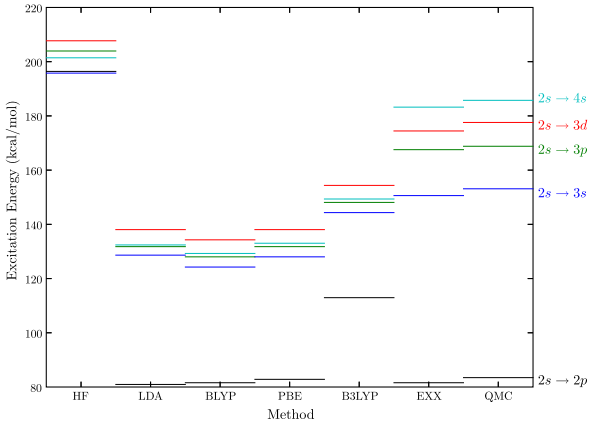

FIG. 7 Comparison of excitation energies for Be atom computed from the KS-DFT eigenvalues using HF, LDA, BLYP, PBE, B3LYP, EXX, and QMC exchange-correlation functionals.

by 0.241–0.273 h, and the B3LYP function has a MAD of 0.15 h. Our EXX values differ by only 0.016 h, nearly a factor of 10 better than the best of the standard DFTs. Moreover, the EXX values differ from the QMC values by almost a constant value of 0.016 h across the entire spectrum, whereas the LDA, GGA, and HF values fluctuate much more about their average deviation. These spectra are shown graphically in Figure 6. This figure illustrates that not only are the EXX values better on the average, but each individual excitation level differs from the corresponding exact value by a near-constant shift.

Table III reports a similar comparison for the Beryllium excitation energies. On the average the HF excitation energies differ by 0.075 h, the LDA, BLYP, and PBE values differ by 0.050-0.070, and the B3LYP values differ by 0.038 h. In contrast, the EXX values differ by only 0.004 h, nearly a factor of ten smaller average difference than the best of the standard DFTs. Moreover, as was seen in He, the EXX excitation energies differ from the QMC values by a constant shift, whereas the LDA, GGA, and HF values fluctuate much more about their averages. These spectra are shown graphically in Figure 7, which again illustrates that not only does the EXX perform better on the average, but that each individual excitation value once again differs from the exact levels by a small, nearly constant shift. We have not yet implemented the EXX-GVB approach for more than two electrons, and thus do not have Be results with this approach.

TABLE IV Comparison of HF, GVB, BLYP, B3LYP, and EXX-GVB energies (in Hartree atomic units) versus H—H (in Angstrom) for H_2 dissociation.

R/Å	HF	GVB	EXX-GVB	BLYP	B3LYP
0.6	-1.1139	-1.1286	-1.1256	-1.1479	-1.1592
0.7	-1.1305	-1.1476	-1.1441	-1.1663	-1.1771
0.8	-1.1284	-1.1484	-1.1447	-1.1663	-1.1765
0.9	-1.1165	-1.1398	-1.1359	-1.1566	-1.1662
1.0	-1.0995	-1.1265	-1.1226	-1.1421	-1.1510
1.2	-1.0594	-1.0956	-1.0915	-1.1079	-1.1150
1.4	-1.0186	-1.0666	-1.0626	-1.0739	-1.0793
1.8	-0.9468	-1.0259	-1.0241	-1.0181	-1.0193
2.2	-0.8904	-1.0067	-1.0061	-0.9795	-0.9764
2.5	-0.8571	-1.0008	-1.0007	-0.9597	-0.9535
3.0	-0.8157	-0.9975	-0.9974	-0.9387	-0.9281
4.0	-0.7702	-0.9965	-0.9965	-0.9214	-0.9050

We believe it significant that for the He and Be spectrum the EXX excitation levels differ from the exact levels by a nearly constant amount. Clearly, the HF description on which our EXX is based omits the correlation energy, and the high accuracy of the excitation levels and the regularity of the error gives hope that simple models for the electron correlation might further reduce the overall error.

4. Results on Dissociation of H_2

Table IV shows energies (in Hartree atomic units) for dissociating H_2 as a function of H—H distance, in Angstroms. For H_2, the OEP-HF results are identical to the HF results, and thus exhibits the same difficulties dissociating that the HF wave function does. BLYP also dissociates to the incorrect limit, although it is closer to the correct result than the HF or EXX results. The GVB and EXX-GVB methods dissociate to nearly the correct limit (the error shown here is predominantly due to finite basis set size effects). The GVB and EXX-GVB methods do not yield exactly the same energies because the orbitals in the EXX-GVB method are slightly constrained by the requirement that they come from a Kohn-Sham equation. This is similar to the energy differences seen between the HF and EXX method (for example, in reference (24)), and was explained particularly well by Kummel and Perdew in reference (10). These data are shown graphically in Figure 8.

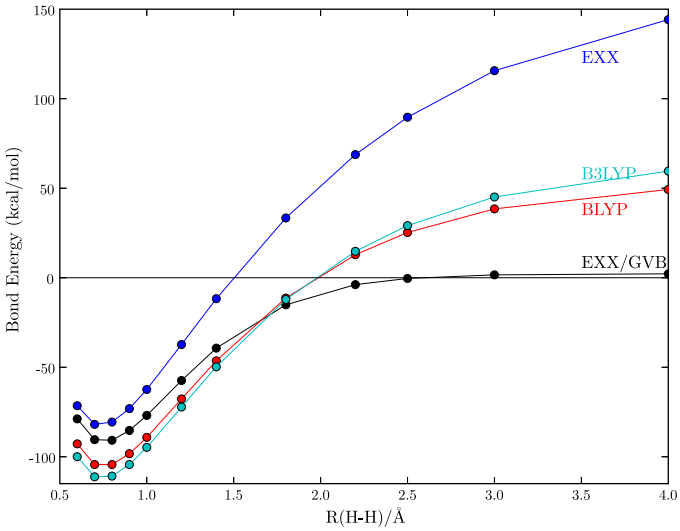

FIG. 8 A comparison of the dissociation of H_2 using HF, GVB, EXX, EXX-GVB, and BLYP methods.

B. Hybrid DFT-QMC Calculations

We are also pursuing the use of quantum Monte Carlo (QMC) techniques as a way of improving the accuracy of DFT calculations. QMC methods take a very different approach to electronic structure than do DFT methods. In DFT one simplifies the exact Hamiltonian to a solvable form; typically this simplification takes the form of reducing the equations to a series of one-particle (or one-pseudoparticle) equations. In contrast, QMC approaches retain the exact form of the Hamiltonian, and use MC techniques to sample the exact form rather than solving it exactly. This enables a solution that is very nearly exact, but at the expense of the slow statistical convergence typical of MC approaches.

DFT KS wave functions are composed of a product of one-particle states

$$\Psi(1,2,\ldots,N_{el}) = \psi_1(r_1)\psi_2(r_2)\cdots\psi_{N_{el}}(r_{N_{el}}). \quad (26)$$

QMC wave functions typically include a correlation term that explicitly describes the electron-electron correlation in terms of the inter-electronic coordinates, in much the same way as the Hylleraas wave function does. The general electron-electron correlation terms are called *Pade-Jastrow* terms and have the form

$$U = \sum_{i<j}^{N_{el}} \frac{a_1 r_{ij} + a_2 r_{ij}^2 + \cdots}{1 + b_1 r_{12} + b_2 r_{12}^2 + \cdots} \quad (27)$$

$$\Psi_{Jastrow} = e^U. \quad (28)$$

The overall QMC many-particle wave function then has the form

$$\Psi_{QMC} = \Psi_{Det}\Psi_{Jastrow} \quad (29)$$

where Ψ_{Det} is an antisymmetrized product wave function made by taking the determinant of a wave function of the form of eq (26).

The KS eigenfunctions that result from DFT calculations generally serve as an accurate guess to the one-particle states in Ψ_{Det}, which reduces the overall problem to one of optimizing the Jastrow parameters in eq (27). This optimization is substantially more complicated than one would normally assume because the objective function to be optimized involves using Monte Carlo techniques to sample the many-body Hamiltonian.

We have collaborated with Cyrus Umrigar from the Cornell Theory Center on the use of DFT calculations as starting points for the increased accuracy of QMC calculations. Earlier CSRF projects have helped fund the development of the CHAMP (Cornell-Holland Abinitio Materials Package) program, in particular the ability to compute materials having periodic boundary conditions, which is essential to Sandia's interests. Our current project simply seeks tighter integration of CHAMP with Sandia's DFT program, in particular the Quest package.

CHAMP has several qualities that make it an ideal choice for our research. The Jastrow optimization uses variance (31) and energy (32) optimization, making it substantially more efficient than other programs in current use. Furthermore, the Monte Carlo sampling is done using very efficient algorithms (33). Finally, the capability for condensed phase systems mentioned above allow materials chemistry properties to be computed.

We have begun work integrating CHAMP with Quest. Two steps need to be achieved for this to occur. First, the wave function, basis set, geometry, and so forth from the DFT calculation needs to be output in a readable form. Much of this functionality has already been incorporated into Quest through various post-processing modules, and this material only needs to be put into the CHAMP format. Secondly, we need to output the pseudopotentials in an appropriate format, and to write suitable routines in CHAMP for these pseudopotentials to be evaluated. The pseudopotentials that CHAMP can handle limited to non-separable forms, and this restriction recommends using Quest above other programs in common use at Sandia (VASP, Socorro, abinit) that require separable pseudopotentials.

C. Future Directions for Increasing the Accuracy of DFT Approaches

The EXX-GVB approach described above is very promising, but in many ways it is unsuitable for simulations of large condensed-phase materials with periodic boundary conditions, because the enumeration of the valence bonds required by the GVB wave function becomes prohibitive. We therefore plan to use the EXX-GVB approach as a benchmark and find other approaches that can achieve similar high-quality results in more tractible form.

There are several obvious choices, including the inclusion of LDA and GGA correlation functionals in our OEP optimizations. There are also other orbital-dependent choices that we believe are promising. The Colle-Salvetti correlation functions (34; 35) is an orbital-dependent functional from which the density-dependent LYP GGA density functional was derived (36). In addition to exploring the use of the LYP functional itself, we believe there are significant advantages to the Colle-Salvetti orbital-dependent functional. The Colle-Salvetti correlation functional is given by (36)

$$E_c = -a \int \frac{1}{1+d\rho(\mathbf{r})^{-1/3}} \times (\rho(\mathbf{r}) + b\rho(\mathbf{r})^{-2/3}[t_{HF}(\mathbf{r}) - 2t_W(\mathbf{r})]e^{-c\rho(\mathbf{r})^{-1/3}})d\mathbf{r}, \quad (30)$$

which includes kinetic-energy terms t_{HF} and t_W that normally only appear in the so-called Meta-GGA DFT approaches, and that may improve the performance of the functional.

Another orbital-dependent choice of a correlation functional comes from many-body perturbation theory, and is particularly attractive given Bartlett and coworkers' success with this approach. We have already explored the use of the second-order Moller-Plesset (MP2) method (37), but did not find this approach to be sufficiently accurate. However, we believe that the closely-related Epstein-Nesbet correlation theory might meet many of our requirements. Second-order Epstein-Nesbet (EN2) theory is given by

$$E_c = \sum_{ab}^{occ} \sum_{rs}^{virt} \frac{(ar|bs)(2(ar|bs) - (br|as))}{E_{ab}^{rs}} \quad (31)$$

where the indices ab run over the occupied orbitals, the indices rs run over unoccupied (virtual) orbitals, the terms $(ar|bs)$ are again in the chemist's notation of eq (22), and the energy E_{ab}^{rs} is the energy of a wave function where two electrons are excited out of orbitals ab into orbitals rs. EN2 has the desirable features of being tractible, variational, size-consistent, and will hopefully lead to proper bond-dissociation.

Finally, or plans for the hybrid DFT-QMC calculation center around implementing the Quest pseudopotentials in the CHAMP program. This capability will enable the QMC approach to also be used as a benchmark in evaluating the performance of OEP correlation functionals, above.

IV. INCREASING THE SPEED OF DFT

A. Density Matrix Purification

DFT consists of self-consistently solving the Kohn-Sham equations. Figure 1 shows a diagram of where time is spent in a DFT calculation. Essentially, there are two time-consuming steps. There is a *Hamiltonian formation step* that forms the Kohn-Sham Hamiltonian operator for a given system for a specified density matrix $\mathbf{D}$, and there is a *Hamiltonian solution step* that solves the eigenproblem associated with that Hamiltonian and forms an updated $\mathbf{D}$ matrix; these two steps are iterated until self-consistency is reached. Formally, the Hamiltonian formation step scales $\mathcal{O}(N^4)$ for a system of size N. However, intelligent Hamiltonian formation steps have reduced the scaling of this step to roughly $\mathcal{O}(N)$ for large systems. The eigenproblem formally scales $\mathcal{O}(N^3)$, and this scaling now dominates the calculation time for large systems.

Figure 2 shows the timings of the Quest DFT program. The red and the blue lines show the timings for Hamiltonian formation during the setup phase and during each iteration, respectively. The green line shows the timings for the solution of the eigenproblem. Beyond roughly 150 atoms, the eigensolve dominates the problem.

But, referring back to Figure 1 we note that one merely uses the eigenvectors from the $\mathbf{H}$ solution step to form a new density matrix.

Our work to date has focused on density matrix purification techniques (38–43) that use the facts that a proper density matrix is idempotent, has a trace equal to the number of electrons, and commutes with $\mathbf{H}$, to refine iteratively, or purify, an initial guess to the density matrix until convergence. Similar approaches include density matrix minimization techniques (44; 45) that apply a nonlinear conjugate gradient minimization to the eigenproblem augmented by penalizing the idempotent constraint, and density matrix expansion techniques (46–50) that expand the Fermi-Dirac distribution in a Chebyshev series to yield the density matrix.

Density matrix purification techniques exploit the fact that a polynomial Φ in $\mathbf{H}$ yields

$$\Phi(\mathbf{H}) = \Phi(\mathbf{U}\mathbf{\Lambda}\mathbf{U}^\mathbf{T}) = \mathbf{U}\Phi(\mathbf{\Lambda})\mathbf{U}^\mathbf{T}, \quad (32)$$

i.e., the polynomial Φ does not change the eigenvectors, and modifies the eigenvalues as $\Phi(\lambda)$. Therefore, $\Phi(\mathbf{H})$ is a density matrix provided that

$$\Phi(\lambda_i) = \begin{cases} 1 & \lambda_i < \mu \\ 0 & \lambda_i > \mu \end{cases}, \quad (33)$$

that is, Φ is an approximation to a step function centered on μ, the chemical potential (defined as the orbital energy below which all orbitals are occupied, and above which all orbitals are unoccupied). In practice the polynomial Φ is formed implicitly by repeatedly multiplying and shifting $\mathbf{D}_j$, the j-th iterative guess to the density matrix, with itself. In the Trace Correcting Canonical Purification Scheme (41; 42), the iterations are determined by

$$\mathbf{D}_{j+1} = \begin{cases} 2\mathbf{D}_j - \mathbf{D}_j^2 & \mathrm{Tr}[\mathbf{D}_j] < N_{el} \\ \mathbf{D}_j^2 & \mathrm{Tr}[\mathbf{D}_j] > N_{el} \end{cases}. \quad (34)$$

Essentially, the purification iterations replace the eigensolve with a series of matrix multiplications. Such

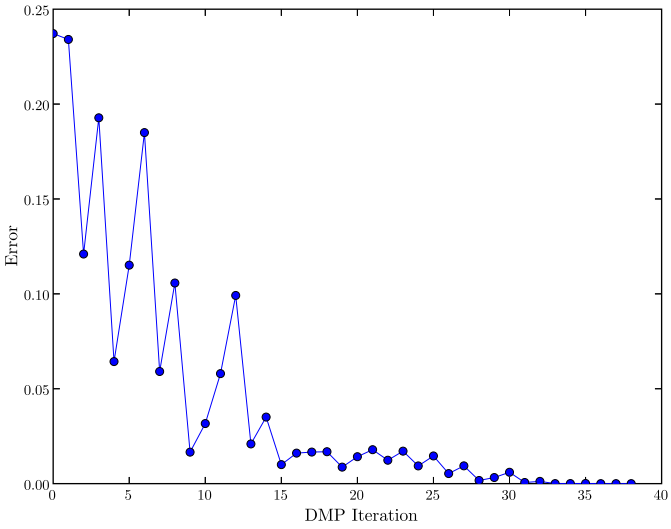

FIG. 9 Convergence of the error during density matrix purification versus purification iteration for an 8-atom Si unit cell.

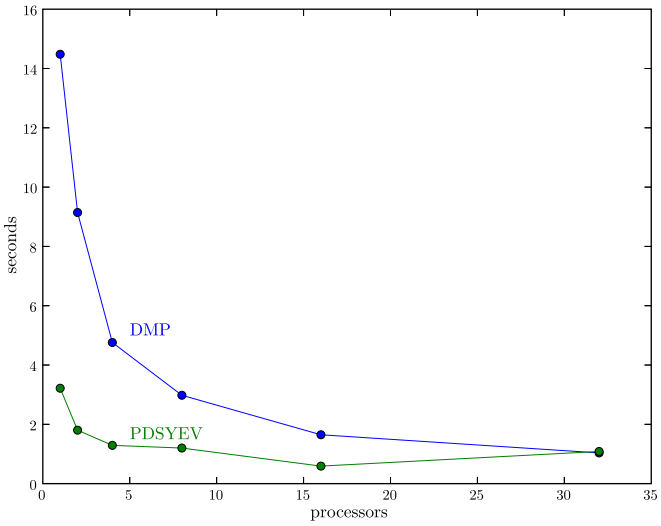

FIG. 10 Comparison of density matrix purification (blue) to the Scalapack routine PDSYEV (green) on the Red Squall parallel computer.

a replacement is in itself nothing unique, since methods like the power method solve the eigenproblem by matrix multiplication. However, what is significant about density matrix purification is that the process is so rapid. Figure 9 shows the rapid convergence of the method for an 8-atom Si unit cell.

Replacing an eigensolve operation with matrix multiplication can potentially provide a significant computational savings because matrix multiplication is an easy operation to take advantage of parallelism and matrix sparsity. Figure 10 shows that this potential savings can in fact be achieved with modest numbers of processors. The purification results (blue) are roughly six times slower than the ScaLAPACK PDSYEV eigensolver (green) for a single processor. However, the parallelization is much faster, so that on 32 processors the speeds are roughly equivalent.

B. Future Directions for Increasing the Speed of DFT

Density matrix purification approaches can outperform dense eigensolvers provided there is parallelism or sparsity that can be taken advantage of. Our investigations to date have used only the parallelism, and have demonstrated that it is possible to perform as well as the dense ScaLAPACK approaches with enough processors. In the remaining time on this project we plan to implement our purification approaches using the Sandia Trilinos (51) framework.

We need to be able to do matrix multiplies of the form

$$\mathbf{C} = \mathbf{AB}. \tag{35}$$

Trilinos uses a compressed row storage (CRS) format for sparse matrices, and, for example, considers the sparsity graphs of matrices $\mathbf{A}$ and $\mathbf{B}$ to determine which elements of $\mathbf{C}$ need be computed.

In DFT software it is often more useful to use block compressed row storage (BCRS) instead of CRS. This is because each atom leads to 10 basis functions, and when two atoms are a long distance apart, the entire 10×10 matrix block containing elements from those two atoms are typically zero. Thus, by using atom-atom distances we can often determine which blocks are zero or nonzero before computing the elements.

Trilinos has the capability to use variable block row storage (VBR), which is ideal for our uses because it can even consider cases where different atoms have different numbers of basis functions (and thus different-sized matrix blocks).

Because Trilinos does not implement matrix multiplication between VBR matrices, we plan to extend the Epetra module to include this functionality. Much of the graph manipulation can be taken directly from the CRS graphs, except now the graphs refer to entire blocks instead of points.

Once we have successfully implemented VBR matrix multiplication, we plan to exploit several other characteristics of our particular problem to achieve better performance. First, we typically require only a matrix square,

$$\mathbf{C} = \mathbf{AA} \tag{36}$$

rather than a general matrix multiply. Furthermore, our matrices are typically real symmetric (or, in the worst case, complex Hermetian), which means we can perform matrix multiplies of the forms

$$\mathbf{C} = \mathbf{A}^T\mathbf{A} \tag{37}$$

$$\mathbf{C} = \mathbf{AA}^T \tag{38}$$

when it allows communications to be reduced. We anticipate that this will result in very high performance for our density matrix purification approach.

V. INCREASING THE RANGE OF DFT

A. Gibbs Ensemble Monte Carlo Approaches With DFT Energies

Elemental metal vapor-liquid coexistence curves (VLCCs) are difficult to obtain via experiment due to the extremely high temperatures required. Consequently data is available only for a handful of the low melting point elements, yet FEM simulations of manufacturing processes like laser welding sensitively depend on this material property information. Gibbs Ensemble Monte Carlo molecular simulations are a proven method for determining VLCCs for small organic molecules and theoretically are an ideal technique for determining atomic metal VLCCs. It is currently unknown whether existing functional forms and parameterizations (force fields) that are fit to reproduce metallic solid and metallic liquid phase data will be appropriate for describing coexistence between a metallic liquid phase and an insulating vapor phase. To meet the need of the manufacturing simulations and cope with the technical challenge, we are pursuing a hybrid DFT-Gibbs MC program to compute the VLCCs using quantum mechanical energies.

B. Gaussian-based Methods for High-Pressure Systems

Gaussian basis sets have many advantages for materials chemistry: they provide a very compact description of the electronic degrees of freedom, because they center their variations on regions in which chemical bonding changes most rapidly. However, one disadvantage is that Gaussian basis sets are not orthogonal: two basis functions χ_μ and χ_ν typically have nonzero overlap:

$$S_{\mu\nu} = \int \chi_\mu(r)\chi_\nu(r)dr. \tag{39}$$

This nonzero overlap means that the Hamiltonian solution step in Figure 1 is, in fact, a generalized eigenproblem

$$\mathbf{HC} = \mathbf{SC}\Lambda \tag{40}$$

rather than the regular eigenproblem shown in eq (5). The generalized eigenproblem is typically solved using a Cholesky decomposition of the overlap matrix $\mathbf{S}$, as part of, for example, the LAPACK routine DSYGV.

However, in high-pressure systems the overlap matrix $\mathbf{S}$ contains small eigenvalues, making the Cholesky decomposition unstable, and producing artifact states in the spectrum of states from eq (40). The solution we have implemented to this problems involves a *Canonical Orthogonalization* procedure that removes small eigenvalues from the overlap matrix, and, with them, the artifact states from the spectrum.

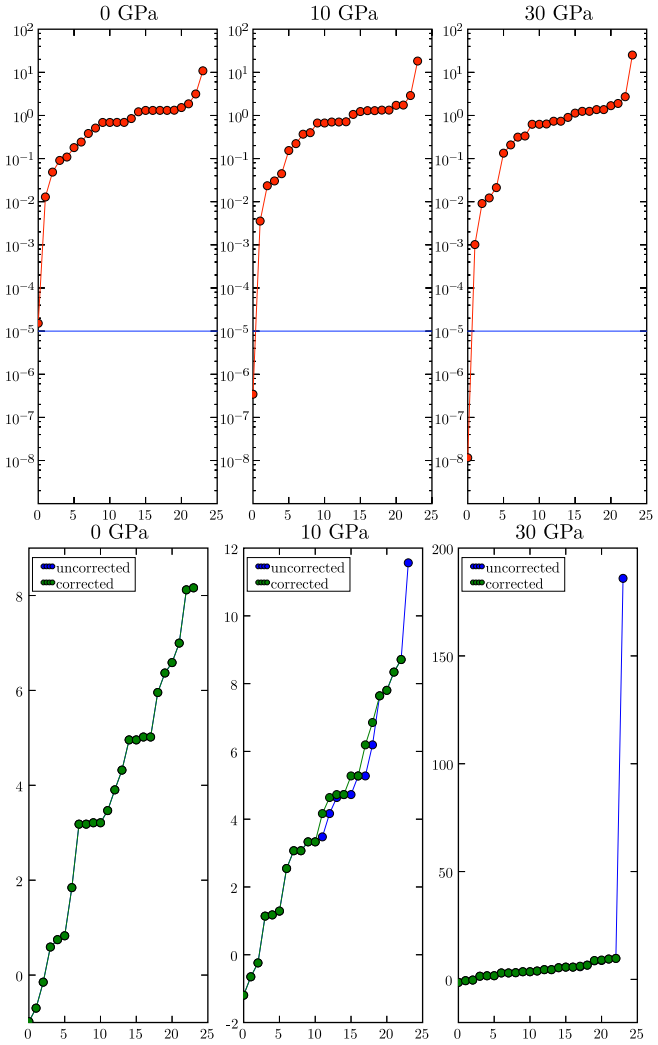

FIG. 11 Results of the Quest high-pressure solver: (top) Overlap spectrum for molecular hydrogen at 0, 10, 30 GPa. The red curve shows the spectrum of the eigenvalue of the $\mathbf{S}$ matrix, and the blue line denotes the cutoff value of 10^-5 used in the Canonical Orthogonalization. (bottom) Corresponding eigenvalues of the $\mathbf{H}$ matrix: There are no artifact states at 0 GPa, as shown by the fact that the corrected curve (green) have the same values as the uncorrected curve (blue). As the pressure rises to 10 GPa these states start to appear, as shown by the difference between the corrected and uncorrected values. By 30 GPa the artifact states completely dominate the higher ends of the spectrum. In all cases, these states are removed by the Canonical Orthogonalization procedure; no change is made to the spectrum by this procedure when artifact states are not present.

We first compute the eigen-decomposition of the overlap matrix

$$\mathbf{SU} = A\mathbf{U} \tag{41}$$

and then compute a transformation matrix $\mathbf{X}$ whose i-th

column is given by

$$\mathbf{X}_i = \mathbf{U}_i/\sqrt{a_i} \tag{42}$$

only for eigenvalues $a_i > \epsilon$, with ϵ a cutoff typically in the range 10^{-7}–10^{-5}. If we define

$$\mathbf{C} = \mathbf{X}\mathbf{C}' \tag{43}$$

$$\mathbf{H}' = \mathbf{X}^T\mathbf{H}\mathbf{T} \tag{44}$$

we can rewrite eq (40) as

$$\mathbf{H}\mathbf{X}\mathbf{C}' = \mathbf{S}\mathbf{X}\mathbf{C}'\Lambda \tag{45}$$

$$\mathbf{H}'\mathbf{C}' = \Lambda\mathbf{C}' \tag{46}$$

which is now transformed into a regular eigenproblem in an orthogonal basis free of the artifact states.

Figure 11 shows the result of this transformation on a molecular hydrogen crystal at differing pressures. The top figures show the spectra of the overlap matrix $\mathbf{S}$, showing that the smallest eigenvalue gets increasingly smaller as the pressure increases. The horizontal line shows the cutoff value used by the Canonical Orthogonalization algorithm. The bottom figures show the corresponding spectra of the Hamiltonian matrices $\mathbf{H}'$. Here, the blue curves show the spectra that contain the artifact states, which are not present at low pressures but begin to dominate the spectrum as the pressure is raised. In contrast, the green curves show the correct spectrum successfully removing the artifact states.

This purification step allows the advantages of the Gaussian basis sets (small overall size coupled with high accuracy because the degrees of freedom are centered at the chemical bond regions) may be used for high pressure simulations.

C. DFT Molecular Dynamics

Classical atomistic force fields replace the DFT equations, which are functions of the electron density, with functions of the atoms alone, effectively "integrating" over the electronic degrees of freedom. In some cases this is an effective substitution. Protein docking, for example, appears to be quite well-described by these classical force fields. However, for more complicated systems, and, in particular, anytime chemistry is occurring, classical force fields are suspect at best.

We have therefore implemented molecular dynamics in the Quest DFT program. This Quest-MD approach uses the quantum mechanical forces computed on each atom to propagate the atoms according to Newtonian dynamics. Such an approach allows complicated behavior like radiation-induced substitution or vacancy relaxation in Silicon, or pressure-induced metallization of molecular hydrogen, to be computed using accurate DFT functionals.

We have implemented several different MD integration schemes, including ones that maintain constant energy and constant temperature. This capability has already been released to the Quest users.

D. Local Coordinate Projection Methods for DFT

When a DFT geometry optimization run needs a large number of iterations to converge, it is often due to the continuing motion of a few atoms in the system. The problem is that the Cartesian coordinate system is sometimes not an efficient representation for some local regions. There are a number of existing schemes for generating internal coordinates for all of the atoms in periodic systems, but they are, in general, costly and complicated. These methods can be mathematically elegant, but in truth, we only need to improve the coordinate system for the problematic atoms, rather than the whole system.

We approach this problem by projecting the forces of selected atoms onto local coordinates that are more appropriate. For example, for bridging atoms (2-fold coordinated) we use a cylindrical coordinate system. The primary advantage is a separation of the stiffer bond mode from rotational motions. In the original cartesian representation, these motions are usually coupled, resulting in slower convergence.

This method only does the projection for selected atoms, or groups of atoms, thus avoiding the unnecssarily complex and costly iterative transformation/back transformation of full periodic projection schemes.

In its current implementation, without full tuning, the cylindrical projection results in a factor of 2 decrease in the number of steps required for convergence.

E. Future Directions for Increasing the Range of DFT

Our focus for the remainder of the CSRF project for increasing the range of DFT will be on the hybrid Gibbs MC-DFT approaches. During the first 18 months of this project we have rewritten parts of Quest so that the program may now be called as a subroutine from the Towhee MC program. Because of the vastly greater resources required for DFT as opposed to classical force fields, we are using the task parallel version of Quest for even the proof-of-concept calculations. We are currently in the process of simulating liquid densities of Si using the hybrid code, which is the first step to computing the VLCCs.

VI. ACKNOWLEDGEMENTS

This work was supported by the Sandia Computer Science Research Foundation. Sandia is a multiprogram laboratory operated by Sandia Corporation, a Lockheed Martin Company, for the United States Department of Energy under Contract No. DE-AC04-94AL85000. This document has been released as SAND 2006-2014P.

References

[1] L. H. Thomas. The calculation of atomic fields. *Proc. Cambridge Philos. Soc.*, 23:542, 1927.

[2] E. Fermi. Un metodo statistice per la determinazione di alcune proprieta dell'atomo. *Rend. Acad., Lincei*, 6:602, 1927.

[3] P. A. M. Dirac. Notes on exchange phenomena in the Thomas atom. *Proc. Cambridge Philos. Soc.*, 26:376, 1930.

[4] Walter Kohn and L. J. Sham. Self-consistent equations including exchange and correlation effects. *Physical Review*, 140:A1133, 1965.

[5] D. M. Ceperley and B. I. Alder. Ground state of the electron gas by a stochastic method. *Physical Review Letters*, 45:566, 1980.

[6] S. H. Vosko, L. Wilk, and M. Nusair. Accurate spin-dependent electron liquid correlation energies for local spin density calculations: A critical analysis. *Canadian Journal of Physics*, 58:1200, 1980.

[7] John C. Slater. A simplification of the Hartree-Fock method. *Physical Review*, 81:385–390, 1951.

[8] Axel D. Becke. Density functional theories in quantum chemistry: Beyond the local density approximation. In Dennis Salahub and Michael Zerner, editors, *The Challenge of d and f electrons: Theory and Computation*, number 394 in ACS Symp. Ser., page 166. American Chemical Society, Washington, 1989.

[9] J. P. Perdew. Density-functional approximation for the correlation energy of the inhomogeneous electron gas. *Physical Review B*, 33:8822, 1986.

[10] J. P. Perdew, K. Burke, and M. Ernzerhof. Generalized gradient approximation made simple. *Physical Review Letters*, 77(18):3865, 1996.

[11] Axel D. Becke. Density functional thermochemistry III: The role of exact exchange. *Journal of Chemical Physics*, 98:5648, 1993.

[12] Axel D. Becke. A new mixing of Hartree-Fock and local density functional theories. *Journal of Chemical Physics*, 98:1372, 1993.

[13] Larry A. Curtiss, Krishnan Raghavachari, Paul C. Redfern, and John A. Pople. Assessment of gaussian-2 and density functional theories for the computation of enthalpies of formation. *Journal of Chemical Physics*, 106(3):1063, 1997.

[14] J. D. Talman and W. F. Shadwick. Optimized effective atomic central potential. *Physical Review A*, 14:36, 1976.

[15] Andreas Görling and Mel Levy. Corrlelation-energy functional and its high-density limit obtained from a coupling-constant perturbation expansion. *Physical Review B*, 47:13105, 1993.

[16] Andreas Görling and Mel Levy. Exact Kohn-Sham scheme based on perturbation theory. *Physical Review A*, 54:196, 1994.

[17] Andreas Görling. Exact treatment of exchange in Kohn-Sham band-structure schemes. *Physical Review B*, 53:7024, 1996.

[18] Martin Stadele, J. A. Majewski, P. Vogl, and Andreas Görling. Exact Kohn-Sham exchange potential in semiconductors. *Physical Review Letters*, 79:2089, 1997.

[19] W. J. Hunt and William A. Goddard, III. Excited states of H_20 using improved virtual orbitals. *Chemical Physics Letters*, 3:414, 1969.

[20] Martin Stadele, J. A. Majewski, P. Vogl, and Andreas Görling. Exact exchange Kohn-Sham formalism applied to semiconductors. *Physical Review Letters*, 79:2089, 1997.

[21] Martin Stadele and Richard M. Martin. Metallization of molecular hydrogen: Predictions from exact-exchange calculations. *Physical Review Letters*, 84:6070, 2000.

[22] Ireneusz Grabowski, So Hirata, Stanislav Ivanov, and Rodney J. Bartlett. Ab initio density functional theory: OEP-MBPT(2). a new orbital-dependent correlation functional. *Journal of Chemical Physics*, 116:4415, 2002.

[23] Rodney J. Bartlett, Victor F. Lotrich, and Igor V. Schweigert. Ab initio density functional theory: The best of both worlds? *Journal of Chemical Physics*, 123:62205, 2005.

[24] Weitao Yang and Qin Wu. Direct method for optimized effective potentials in density functional theory. *Physical Review Letters*, 89:143002, 2002.

[25] Frank W. Bobrowicz and William A. Goddard, III. The self-consistent field equations for generalized valence bond and open-shell Hartree-Fock wave functions. In Henry F. Schaefer, III, editor, *Methods of Electronic Structure Theory*, volume 3 of *Modern Theoretical Chemistry*, page 79. Plenum Press, New York, 1977.

[26] Atilla Szabo and Neil S. Ostlund. *Modern Quantum Chemistry*. McGraw-Hill, New York, 1982.

[27] A. Savin, Cyrus J. Umrigar, and Xavier Gonze. Relationship of Kohn-Sham eigenvalues to excitation energies. *Chemical Physics Letters*, 288:391, 1998.

[28] S. Bashkin and J. D. Stoner. *Atomic Energy Levels and Groatian Diagrams*. North-Holland, Amsterdam, 1975.

[29] G. W. F. Drake and Z. C. Yan. Variational eigenvalues for the S-states of Helium. *Chemical Physics Letters*, 229:486, 1994.

[30] G. W. F. Drake. In J. Seminario, editor, *Casimir Forces: Theory and Recent Experiments on Atomic Systems*. Elsevier, Amsterdam, 1997.

[31] C. J. Umrigar, K. G. Wilson, and J. W. Wilkins. Optimized trial wave functions for quantum Monte Carlo calculations. *Physical Review Letters*, 60:1719, 1988.

[32] C. J. Umrigar and C. Filippi. Energy and variance optimization of many-body wave functions. *Physical Review Letters*, 94:150201, 2005.

[33] C. J. Umrigar, M. P. Nightingale, and K. J. Runge. A diffusion quantum Monte Carlo algorithm with small time step errors. *Journal of Chemical Physics*, 99:2865, 1993.

[34] Renato Colle and Oriano Salvetti. Approximate calculation of the correlation energy for the closed shells. *Theoretica Chimica Acta*, 37:329, 1975.

[35] Renato Colle and Oriano Salvetti. An iterative method for the direct calculation of the total electronic energy and correlated wave function through the use of the colle-salvetti correlation potential. *Journal of Chemical Physics*, 94:1306, 1990.

[36] C. Lee, W. Yang, and R. G. Parr. Development of the colle-salvetti correlation energy formula into a functional of the electron density. *Physical Review B*, 37:785, 1988.

[37] C. Møller and M. S. Plesset. *Physical Review*, 46:618, 1934.

[38] R. McWeeny. Some recent advances in density matrix theory. *Reviews of Modern Physics*, 32:335, 1960.

[39] A. H. R. Palser and D. E. Manolopoulos. Canonical purification of the density matrix in electron structure the-

ory. *Physical Review Letters*, 58, 1998.

[40] G. Beylkin, N Coult, and M. J. Mohlenkamp. Fast spectral projection algorithms for density-matrix computations. *Journal of Computational Physics*, 152:32, 1999.

[41] A. M. N. Niklasson. Expansion algorithm for the density matrix. *Physical Review B*, 66:155115, 2002.

[42] A. M. N. Niklasson. Implicit purification for temperature dependent density matrices. *Physical Review B*, 68:233104, 2003.

[43] A. M. N. Niklasson, C. J. Tymczak, and M. Challacombe. Trace resetting density matrix purification in o(n) self consistent field theory. *Journal of Chemical Physics*, 118:8611, 2003.

[44] X.-P. Li, R. W. Nunes, and David Vanderbilt. Density matrix electronic structure method with linear system size scaling. *Physical Review B*, 47:10891, 1993.

[45] M. S. Daw. Model for energetics of solids based on the density matrix. *Physical Review B*, 47:10895, 1993.

[46] J. M. Millan and G. E. Scuseria. Linear scaling conjugate gradient density matrix search as an alternative to diagonalization for first principles electronic structure calculations. *Journal of Chemical Physics*, 106:5569, 1997.

[47] K. R. Bates, A. D. Daniels, and G. E. Scuseria. Comparison of conjugate gradient density matrix search and chebyshev expansion methods for avoiding diagonalization in large scale electronic structure calculations. *Journal of Chemical Physics*, 109:3308, 1998.

[48] A. D. Daniels and G. E. Scuseria. What is the best alternative to diagonalization of the hamiltonian in large scale semiempirical calculations? *Journal of Chemical Physics*, 110:1321, 1999.

[49] K. Nemeth and G. E. Scuseria. Linear scaling density matrix search based on sign matrices. *Journal of Chemical Physics*, 113(15):6035, 2000.

[50] S. Goedecker and G. E. Scuseria. Linear scaling electronic structure methods in chemistry and physics. *Computing in Science and Engineering*, 5:14, 2003.

[51] Michael A. Heroux, Roscoe A. Bartlett, Vicki E. Howle, Robert J. Hoekstra, Jonathan J. Hu, Tamara G. Kolda, Richard B. Lehoucq, Kevin R. Long, Roger P. Pawlowski, Eric T. Phipps, Andrew G. Salinger, Heidi K. Thornquist, Ray S. Tuminaro, James M. Willenbring, Alan Williams, and Kendall S. Stanley. An overview of the Trilinos project. *ACM Transactions on Mathematical Software*, 31:397, 2005.